

CBO MEMORANDUM

ESTIMATING THE COSTS OF
ONE-SIDED BETS:
HOW CBO ANALYZES PROPOSALS
WITH ASYMMETRIC UNCERTAINTIES

October 1999

20020313 091



CONGRESSIONAL BUDGET OFFICE
SECOND AND D STREETS, S.W.
WASHINGTON, D.C. 20515

A&T02-06-0977

INTERNET DOCUMENT INFORMATION FORM

A . Report Title: Estimating the Costs of One-Sided Bets: How CBO Analyzes Proposals with Asymmetric Uncertainties

B. DATE Report Downloaded From the Internet: 03/12/02

C. Report's Point of Contact: (Name, Organization, Address, Office Symbol, & Ph #): Congressional Budget Office
Second and D Streets, SW
Washington, DC 20515

D. Currently Applicable Classification Level: Unclassified

E. Distribution Statement A: Approved for Public Release

F. The foregoing information was compiled and provided by:
DTIC-OCA, Initials: __VM__ **Preparation Date** 03/12/02

The foregoing information should exactly correspond to the Title, Report Number, and the Date on the accompanying report document. If there are mismatches, or other questions, contact the above OCA Representative for resolution.

In providing the Congress with formal and informal estimates of the costs of legislative proposals, the Congressional Budget Office (CBO) has determined that a small fraction of such proposals cannot be appropriately analyzed by using traditional methods. In keeping with CBO's practice of disclosing its analytical assumptions and methods, this memorandum outlines an alternative method known as probabilistic estimating that the agency has developed and used occasionally since 1995.

Perry Beider of CBO's Natural Resources and Commerce Division wrote the memorandum under the direction of Jan Paul Acton, Roger Hitchner, and David Moore, with assistance from Robert Arnold, David Hull, Craig Jagger, Deborah Kalcevic, Mary Maginniss, and Judith Ruud. Leah Mazade edited the memorandum, Chris Spoor proofread it, and Angela Z. McCollough prepared it for publication. Laurie Brown prepared the electronic versions for CBO's World Wide Web site (www.cbo.gov).

CONTENTS

INTRODUCTION	1
COST ESTIMATING AND "ONE-SIDED BETS"	2
OVERVIEW OF THE PROBABILISTIC APPROACH	5
EXAMPLES OF THE PROBABILISTIC APPROACH	9
Case Study 1: Setting Agricultural Price Support Rates	9
Case Study 2: Changing the Formula for Interest Rates on Student Loans	14
Case Study 3: Paying for Deposit Insurance Losses	19
CONCLUSION	22

INTRODUCTION

One of the primary responsibilities of the Congressional Budget Office (CBO) is to analyze the budgetary effects of legislative proposals. CBO issues formal estimates for every bill reported out of committee as well as every bill passed by the Congress and cleared for the President's signature that affects direct spending or revenues. (Direct spending is spending not subject to annual appropriation.) Earlier in the legislative process, CBO analysts may provide Congressional staff with informal estimates as bills are developed and refined. The agency's estimates play a significant role in the budget procedures of the House of Representatives and the Senate; accordingly, CBO strives to base each estimate on careful analysis and the best available information.¹

In keeping with that goal, CBO has developed a method of analysis known as probabilistic estimating for use in certain rare circumstances in which traditional methods would yield misleading results. This memorandum reviews the traditional approach and its limitations for analyzing bills that involve asymmetric budgetary uncertainties, outlines the basic approach used in probabilistic estimating, and provides brief case studies that illustrate how CBO has used that approach in analyzing legislative proposals.

1. For legislative provisions involving changes in taxes, section 201 of the Congressional Budget and Impoundment Control Act of 1974 requires CBO to use the revenue estimates of the Joint Committee on Taxation.

COST ESTIMATING AND "ONE-SIDED BETS"

For each cost estimate, CBO compares federal spending and revenues under the proposal with the corresponding levels expected under current law. Each estimate projects the incremental costs or savings for each fiscal year of a specified estimation period, typically five or 10 years. For consistency, CBO generally uses the same assumptions in its estimating—often called baseline assumptions—throughout a budget cycle.

In preparing estimates, CBO usually contends with many sources of uncertainty about the future, such as the way the executive branch will implement a bill once it becomes law, the prices of particular goods and services, the performance of the national economy, the weather, technological change, and events in foreign countries. CBO analysts may try to narrow the range of uncertainty surrounding the effects of a bill—for example, by consulting with an implementing agency to gauge the likely scope, purpose, and timing of spending; by analyzing relevant data (on prices, participation in existing programs, economic cycles, and so forth); and by soliciting the views of outside experts. Such efforts reduce but do not eliminate the uncertainty. Thus, in most cases, the ultimate budgetary effect of a bill can never be known with precision.

Nonetheless, the Congressional budget process calls for point estimates. All of the major components of the process—the concurrent budget resolution, reconciliation instructions, appropriations, and enforcement mechanisms such as discretionary spending caps and sequestration—involve specific dollar amounts for specific years. As a result, most of CBO's estimates are presented as point estimates even when the underlying analysis suggests a range of possible outcomes.

Identifying an appropriate point estimate does not usually require the CBO analyst to undertake a detailed analysis of the uncertainty regarding a proposal's budgetary effect. Often, for example, the distribution of possible effects is approximately symmetric around a single, most likely result. In such a case, the analyst can reasonably take a "deterministic" approach, focusing on the scenario underlying the most likely outcome without attempting to quantify the implications of the less likely outcomes on either side of it. Although the resulting estimate will rarely turn out to be exactly correct, the symmetry of the uncertainty implies that the estimate is about equally likely to be too low or too high. Symmetry also suggests that the dollar sum of positive and negative errors from a large number of such estimates will be close to zero.²

2. An analysis of appropriation bills from 1993 through 1997 found that the sum of CBO's estimates of outlays was within 0.1 percent of the actual total during that period, suggesting that the individual estimates were subject to random uncertainty but not to systematic error. See Congressional Budget Office, *An Analysis of CBO's Outlay Estimates for Appropriation Bills, Fiscal Years 1993-1997*, CBO Memorandum (October 1998). Note, however, that appropriation bills generally present a narrower estimating task than do the kinds of bills discussed here: they typically specify levels of budget authority, leaving only the timing of the resulting outlays to be estimated.

The possible costs of a legislative proposal, however, are not always distributed symmetrically around a single, most likely outcome. In some cases, asymmetry is inherent in the uncertain future circumstances: for example, the possibility of massive destruction (and resulting costs) from a catastrophic earthquake or hurricane is not balanced by an equal possibility of natural forces suddenly constructing or repairing a comparable amount of property. In other cases, asymmetry results from particular legislative provisions, such as those with spending floors, caps, or trigger levels. A bill to provide additional benefits when the rate of unemployment goes above 8 percent, for instance, yields increased spending if the actual rate turns out to be 9 percent, but it does not yield symmetric reductions in spending if the rate proves to be 7 percent. As discussed below, analyzing budgetary effects in the presence of asymmetric uncertainty may require going beyond evaluation of a single scenario to consideration of a range of possible outcomes.

A legislative proposal containing a threshold for spending or receipts (for instance, a floor or a cap) can be viewed as a one-sided bet. First, it is a bet because it makes spending or revenues conditional on the outcome of one or more quantifiable but uncertain variables, such as interest rates, unemployment rates, market prices, or offsetting federal receipts. Second, it is one-sided because the budgetary effects (spending or revenues) occur only if the threshold is reached—there are no corresponding opposite effects if the threshold is not reached.

Since 1995, when CBO first began to use probabilistic methods to analyze one-sided bets, the agency has encountered dozens of such bills and proposals. They have involved a broad spectrum of programs covering agricultural commodities, student loans, rural electricity, banking, Medicare, welfare, and the use of receipts by the Departments of Commerce, Defense, Energy, Interior, and Veterans Affairs. Both authorization and appropriation bills have included one-sided bets. At least seven have been enacted into law since 1995, two of which CBO estimated would save the government money.³

OVERVIEW OF THE PROBABILISTIC APPROACH

CBO developed probabilistic methods when it became clear that the standard, deterministic approach focusing on baseline assumptions yields inappropriate results in certain circumstances. For example, many one-sided bets set thresholds for spending at levels that are almost—but not quite—reached under CBO's baseline assumptions. Using the standard approach, CBO would estimate the cost of such proposals as zero. In light of the uncertainty surrounding the baseline, however,

3. The Federal Agriculture Improvement and Reform Act of 1996 included two one-sided bets (on rates of crop price supports, discussed below, and farm acreage set-asides) as did the Balanced Budget Act of 1997 (involving interest payable by states that have received advances from the Unemployment Trust Fund and Medicare capitation payments to health maintenance organizations). The Higher Education Amendments of 1998 (discussed below) and the 1998 appropriations for the Departments of the Interior and Veterans Affairs contained other one-sided bets. The two bets that CBO estimated would reduce federal spending were on crop price supports and the Unemployment Trust Fund.

that result would be misleading: it neglects the possibility that actual circumstances could deviate from the baseline and the threshold might be reached, triggering spending.

Using the probabilistic approach, the analyst estimates the expected value of a proposal's budgetary effect—that is, the weighted average of the effects associated with all possible sets of circumstances, taking account of their respective probabilities. Hence, the approach is also known as expected-value estimating. In principle, using weights equal to the probabilities of the various outcomes can yield budget estimates that are unbiased: if the assumptions of a particular analysis are accurate, then the estimate equals the average of the costs that would be observed if the proposal could be implemented over and over again, as if in a cosmic laboratory, with each case differing only in the way the uncertain factors turn out. A more practical reason to strive for statistical unbiasedness is that estimating errors caused by the inability to forecast uncertain factors should average out to zero over a large number of proposals, provided that the sources of uncertainty affecting the proposals are reasonably independent of each other.

To see the importance of the probability distribution in the expected-value approach, consider the problem of estimating the higher of the two numbers showing when two dice are rolled. That number could be any integer from 1 to 6; thus, the simple unweighted average of the possible outcomes is $3.5 ([1+2+3+4+5+6]/6)$.

However, the chances are greater that one die or the other shows a high number than that both dice show low numbers. Consequently, the probability distribution skews toward the larger numbers, and the expected value is about 4.5.⁴

Identifying the appropriate probability weights to use in calculating the expected value is the major challenge posed by probabilistic estimating. The methods an analyst applies to that task may vary, depending on the proposal and the availability of data. In some cases, historical data may readily suggest a suitable distribution; in others, the analyst may rely more on professional judgments regarding the range of possibilities and the general shape of the distribution (for example, whether it is bounded at specific lower and upper limits or unbounded, whether it is symmetric or skewed, and whether it is flat or single- or multiple-peaked). Thus, CBO may use different methods to analyze two proposals to allow government agencies to spend any receipts they collect above some threshold if one proposal involves visitor receipts collected by the National Park Service, for which years of data exist, and the other applies to receipts from auctions conducted by the Federal Communications Commission (FCC) for licenses for a radically new type of communication service.

4. Assuming the dice are not loaded, the chances are 1 in 36 that both turn up a 1, 3 in 36 that the higher number is a 2, 5 in 36 that it is a 3, and so on. Thus, the expected value is $([1 \times 1] + [3 \times 2] + [5 \times 3] + [7 \times 4] + [9 \times 5] + [11 \times 6])/36 = 161/36$, or approximately 4.472.

Another way in which efforts to identify probability weights for an expected-value estimate differ is that some efforts focus on a close proxy of the budgetary effect itself—as in the previous examples, in which the Park Service and FCC would be authorized to spend all receipts above the thresholds. Other efforts seek a distribution for one or more underlying factors, such as interest rates, crop prices, and losses from natural disasters, that influence the budgetary effect. In those cases, the analyst will also typically need to estimate the budgetary results associated with each relevant outcome (for example, interest rates of 6.5 percent, 6.6 percent, and so on, up to 9.5 percent), thereby converting the probability distribution of the underlying factor or factors into a distribution of the possible budgetary effects. Conceptually, that additional step is the same as it is in the traditional deterministic approach.

For most legislative proposals, the relevant probability distributions are so nearly symmetric, or so hard to distinguish as asymmetric with the limited data available, that using the probabilistic approach would yield little or no improvement in the accuracy of the estimate.⁵ Moreover, the extra complexity of the approach

5. The equivalence of deterministic and probabilistic methods in cases of symmetric uncertainty can be illustrated by CBO's estimates of the federal costs of the so-called "Winstar" lawsuits brought by savings and loan institutions (S&Ls). (The lawsuits focus on the issue of supervisory goodwill, alleging that the government breached agreements it had reached with certain S&Ls regarding the accounting treatment of intangible assets acquired by buying or merging with other institutions.) In its baseline for fiscal year 1997, CBO included \$9 billion over six years for estimated payouts on the lawsuits, representing a middle ground between available upper-bound estimates of \$18 billion to \$20 billion and a conceivable but unlikely lower-bound cost of zero. Because CBO analysts believed that the probabilities of different costs were roughly symmetric within that range, the same \$9 billion estimate could have been derived from either a deterministic approach (focusing on a most-likely scenario in which the government loses about half—\$9 billion worth—of the cases) or a probabilistic approach (using equal probabilities on each side of \$9 billion).

introduces additional possibilities for error and makes the results less transparent. Consequently, CBO continues to rely primarily on the deterministic approach, using probabilistic methods only when they are essential to produce meaningful estimates, as in the cases described below.

EXAMPLES OF THE PROBABILISTIC APPROACH

Probabilistic estimates represent only a small fraction of the thousands of formal and informal estimates that CBO analysts have done since 1995, yet in absolute numbers, they probably total more than 100 as of August 1999. This section offers three brief case studies to illustrate the range of federal programs to which CBO has applied the approach. The three cases—involving crop price supports, student loans, and deposit insurance—also show some of the different methods analysts have used to identify the relevant probability distributions.

Case Study 1: Setting Agricultural Price Support Rates

Between July 1995 and April 1996, CBO analyzed a variety of options for higher and lower rates of price support in the programs providing marketing-assistance loans for major crops—wheat, feed grains, soybeans, cotton, and rice. Under those

programs, farmers can use their crops as collateral for loans from the government. The size of a loan equals the quantity of the crop serving as collateral multiplied by the price support, or loan, rate, which is a unit value (such as dollars per bushel) set by the government.

The loan rate is known as a price support rate because farmers have the option of repaying their loans at any time at the lower of the following: either the outstanding loan principal plus interest or the prevailing market value (as determined daily or weekly by the government) of the quantity of the collateral crop. Similarly, at the end of the loan period, farmers may choose to forfeit their crop in lieu of repaying the amount they owe. Those options effectively guarantee that participating farmers will get at least the price support rate for the crops they borrow against. The benefits received by farmers who take advantage of a positive difference between the price support rate and the prevailing market price are known as marketing loan gains.⁶

During the negotiations over reconciliation in 1995 and over the farm bill in 1996, Members of Congress generated dozens of proposals to change the rules for setting price support rates. The formulas in place since the 1985 farm bill had set the rates for most crops at between 80 percent and 85 percent of a moving average

6. In some cases, when the market price of a crop falls below the preestablished loan rate, farmers can earn a marketing loan gain by placing a loan and then immediately redeeming it. Farmers in that situation who agree to forgo eligibility for a loan are given loan deficiency payments—which for present purposes can be viewed as another form of marketing loan gain.

of past market prices. Some of the proposals sought to lower the rates to 70 percent or 75 percent, others were aimed at raising the rates to percentages in the upper 80s or 90s, and still others sought to keep the existing percentages but cap the dollar levels of the rates (for example, at their 1995 levels).

Those proposals had substantial consequences for the federal budget, potentially increasing or decreasing by billions of dollars annually the marketing loan gains flowing from the government to farmers. Moreover, the consequences were highly and asymmetrically uncertain: under the marketing loan programs, low crop prices yield large costs to the government (with farmers paying their loans back at the market price rather than the loan rate), but high prices provide no offsetting gains (farmers simply repay their loans at the original loan rate, leaving the "up" side for the government capped at zero).

Given that asymmetry, conventional estimating would have given the Congress misleading signals about the costs of alternative proposals. In particular, because CBO's baseline projections assumed relatively high market prices for most crops during the estimation period (1996 to 2002), applying conventional methods to any proposal to increase the rates of price support toward but not above those baseline prices would have yielded a budget estimate of zero. However, such methods ignore the increased chance that actual crop prices will fall below the higher price supports and hence the increased chance (and greater potential size) of farm-

ers' marketing loan gains. Conversely, using conventional methods to estimate costs would have credited little or no federal savings to proposals to reduce or cap loan rates (since the baseline price assumptions implied very low marketing loan gains through 2002), despite the government's reduced exposure to the losses that could occur if actual crop prices turned out to be lower than anticipated.⁷

Ideally, a probabilistic analysis of proposals affecting the rates of price supports would be based on a detailed model of the potential variability of the supply of and demand for each commodity, as affected by weather, economic conditions, commodity inventories, and the provisions of a particular legislative proposal. CBO analysts determined, however, that developing and using such a detailed model to analyze the many proposals developed by the Congress would not be practical.

Accordingly, CBO focused on crop prices—the central uncertainty affecting the outcome for the budget—and assumed that the uncertainty surrounding the price of each crop could be represented by a single probability distribution of deviations from the baseline projection of future prices. To obtain the probability distribution for each crop, CBO first analyzed data on deviations from historical trends for the

7. The inaccuracies resulting from conventional estimating methods would have been different in different circumstances. For example, if the crop prices assumed in the baseline had been at or just slightly below the existing price support rates rather than well above them, conventional methods would have assumed that a small increase in rates would shift the probability of marketing loan gains all the way from zero to 100 percent, thus tending to overstate the incremental costs.

crop's yields to determine the probability of supply shocks of different sizes, then estimated the deviation in price corresponding to each size of shock by applying the demand elasticity used in CBO's estimating model for that crop. To simplify subsequent calculations using the price distributions, analysts grouped the data into discrete price points—adding up, for example, the estimated probabilities of all deviations resulting in corn prices of between \$2.05 and \$2.14 per bushel and assigning that combined probability to the \$2.10 price point.⁸

The development of that relatively simple model of uncertainty in crop prices, which made it possible for CBO analysts to quickly estimate the expected-value costs of proposals to change price support rates, cast the 1995-1996 debate over the farm bill in a new light. The model allowed analysts to estimate not only the expected costs associated with increases in the rates but also the average savings to be expected from reductions. Partly as a result, the option that was finally enacted capped the rates for wheat, corn, cotton, and rice at their 1995 levels. CBO credited the provisions of the bill involving the loan programs with saving \$1.44 billion between 1996 and 2002—two-thirds of the \$2.14 billion savings estimated for the bill as a whole.

8. Because the relatively short record of available data yielded distributions with irregular shapes, analysts also smoothed the distributions to reflect the assumption that larger deviations from the baseline are less likely than smaller deviations.

As it turned out, crop prices in the 1998 market season were lower than those assumed in CBO's 1996 baseline projections, and prices early in 1999 are lower as well, yielding the kinds of large marketing loan gains that the probabilistic estimates of changes in price support rates took into account. Of course, CBO analysts did not know in 1996 that future crop prices would be below those in the baseline, only that they might be—and that the uncertainty in those prices had asymmetric implications for federal spending. Subsequent events have sharply illustrated the importance of analyzing possibilities that differ from the baseline when they have such asymmetric budgetary effects.

Case Study 2: Changing the Formula for Interest Rates on Student Loans

During the Congress's consideration of legislation that became the Higher Education Amendments of 1998, CBO analyzed dozens of proposals affecting the federal costs of student loans for postsecondary education. The student-loan program covers two different types of loans: direct loans from the federal government to students and parents as well as loans issued by private lenders with repayment guaranteed by the government.

Both types of loans carry variable interest rates that are adjusted annually in July according to formulas specified in law. The rates that borrowers must pay,

however, are limited by statutory caps, and those caps have important implications for the federal budget. In the case of direct loans, the caps limit borrowers' repayments, thus raising the overall federal cost of the loans as measured under the estimating rules of credit reform.⁹ For guaranteed loans, whenever the formula-based rates exceed the caps, the government makes so-called special-allowance payments to private lenders to cover the difference. Given the asymmetry in the budgetary effects associated with the caps, estimating the cost of proposals that affect the likelihood of interest rates reaching the caps requires probabilistic methods.

Under an earlier law, new formulas for the interest rates were to take effect on July 1, 1998. The new formulas specified a rate equal to the rate on Treasury bonds of comparable maturity—which CBO interpreted to mean 10-year bonds—plus 1.0 percentage point for loans to students and 2.1 percentage points for loans to parents of students. As in previous years, however, the rates actually paid by borrowers were to be capped at 8.25 percent for student loans and 9.0 percent for parent loans.

The 105th Congress explored many proposals to change the interest rate formulas or the caps, and each proposal would have affected the likelihood that future rates paid by borrowers would reach the caps. The approach reported out of committee in H.R. 6 and ultimately enacted into law left the caps unchanged but revised

9. Under credit reform rules, the budgetary cost of loans made in a given year is the difference between the amount of the loans and the discounted present value of future repayments. Viewed another way, it is the amount of money required to be set aside in an interest-bearing account to finance the expected levels of future defaults and interest subsidies.

the formulas. Under the new law, student loans issued beginning July 1, 1998, carry an annually adjusted interest rate equal to the rate on the three-month Treasury bill plus an adjustment factor of 1.7 percentage points while the borrower is in school or in a grace or deferment period and 2.3 percentage points otherwise. However, lenders of guaranteed student loans get 0.5 percentage points (or 50 basis points) more in interest, with the federal government paying the difference. Loans to parents carry a variable interest rate equal to the three-month Treasury bill rate plus 3.1 percentage points, with no supplemental payment to private lenders.

The new interest rate formulas affect the federal budget in several ways. In one respect, the new formulas lower federal costs: because they specify rates that tend to be higher than those anticipated under previous law, the government gets larger payments on the direct loans. Conversely, however, those higher rates also boost federal costs for in-school subsidies and other payments that are tied to the interest rates. Moreover, by raising the rates specified by the formulas and shifting their basis from the 10-year bond to the more volatile three-month bill while leaving the rate caps unchanged, the new law makes it more likely that the formula-based rates will exceed the caps. That effect serves to limit the increase in repayments on direct loans and bolsters the likelihood and cost of special allowances on guaranteed loans. Finally, the interest supplement of 50 basis points to private lenders directly increases federal costs on all guaranteed loans to students.

In its analysis of H.R. 6 as reported by the House Committee on Education and the Workforce, CBO estimated the cost of the interest rate provisions at about \$8.1 billion from 1998 through 2008 on a loan volume of roughly \$430 billion.¹⁰ That estimate is the discounted present value of the stream of losses expected over the lifetimes—in some cases, as long as 25 or 30 years—of the loans issued during those 11 years. Had CBO ignored the possible budgetary significance of the rate caps in comparing the bill with current law—that is, had it used traditional deterministic methods—it would have estimated \$2.3 billion as the cost of those provisions. The remaining \$5.8 billion represents the expected incremental cost (in special allowances and forgone repayments) of making it more likely that the caps will be reached.

In developing that estimate, CBO's analysis focused on the uncertainty in the formula-based interest rates as the key factor, just as the analysis of changes to the agricultural price support rates concentrated on the uncertainty surrounding the prices of crops. Again, many other uncertain factors are relevant to estimating federal costs for the student-loan programs (for example, the number, dollar volume, and rate of default of loans of each type), but those factors are not associated with asymmetry in budgetary outcomes.

10. CBO provided similar estimates for later versions of the bill as it moved through the Congress. The discussion here focuses on the April 17, 1998, estimate for the House bill as reported to avoid describing the effects of subsequent additional provisions (such as phase-in rules) that interact with the interest rate provisions.

As in case study 1, CBO analysts used historical data to derive the relevant probability distribution—in particular, quarterly data from 1955 to 1997 on interest rates for the 10-year Treasury bond and three-month Treasury bill, on inflation (as measured by the consumer price index), and on the rate of unemployment (as a measure of the strength of the economy). They used the data differently than in the first case, however, in light of the correlation of interest rates from one year to the next. Interest rates tend to move more gradually over time—metaphorically, to have more "inertia"—compared with wheat prices, which may be high one year and low the next (depending on weather conditions, for instance).

Accordingly, CBO used the data to estimate not a static, or "snapshot," probability distribution for the interest rate on three-month Treasury bills but a dynamic model. In the model, the rate at a given point in time is determined not only by contemporaneous random factors but also by the preceding three quarters' inflation, unemployment, and interest rates. After estimating the model and calculating the variance of the residual random component, CBO analysts created 1,000 hypothetical future trajectories of interest rates. Each trajectory started at the same point and evolved according to the model, but each experienced different quarterly random shocks, which were drawn from a normal distribution with a mean (average) of zero and variance equal to that estimated from the data. Finally, analysts used the distribution of the 1,000 simulated trajectories to calculate how often and by how much

the rate for the three-month Treasury bill might be expected to exceed the interest caps during the estimation period.

Case Study 3: Paying for Deposit Insurance Losses

In 1995, the Administration proposed several reforms affecting the Savings Association Insurance Fund (SAIF), which insures deposits at member savings and loan (or thrift) institutions. One reform aimed to prevent the SAIF from shouldering large losses from failed S&Ls (and having to assess heavy additional premiums on the remaining thrifts). The reform would have allowed the SAIF to tap the Treasury to pay for losses above \$500 million in any one year.

At the time of the proposal, CBO's baseline estimate of annual losses from failed thrift institutions was \$450 million—below the \$500 million threshold. Focusing on the baseline figure as the relevant one, conventional deterministic methods would thus have assessed the proposal as having no cost. Again, however, significant and asymmetric uncertainty made that answer inappropriate. The uncertainty was significant because of a substantial probability that actual losses in a given year would exceed the baseline estimate by enough (that is, by more than \$50 million) to trigger federal payments to the SAIF. Moreover, the uncertainty was

entirely one-sided in its budgetary implications—the SAIF would not be required to reimburse the Treasury in years when thrift losses were low.

After reviewing the available data, CBO determined that they did not shed useful light on future annual losses by thrift institutions and that professional judgment would provide the firmest foundation on which to base a probability distribution. Data on how much the former Federal Savings and Loan Insurance Corporation paid to close failed thrifts revealed more about the levels of funding available to the corporation than about the thrift failures themselves. And analysts concluded that more meaningful data—estimates of true losses in the thrift industry and data on bank failures over 20 years—were too old to reflect major changes in legislation, technology, and market conditions affecting the operation and regulation of the thrift industry in the 1990s.

In specifying an appropriate probability distribution, CBO analysts focused on the losses themselves rather than on interest rates, rates of growth for the economy, or other underlying factors. They sought a distribution that met two criteria: it ruled out negative or arbitrarily large losses; and it skewed to the left, assigning a greater likelihood to losses below the mean. Analysts first determined that those criteria could be satisfied by certain beta distributions (a family of distributions, like the better-known normal family, described by a common mathematical formula but distinguished from each other by the values of certain parameters). They then chose

parameters to specify a distribution that bound the possible annual losses between zero and \$4 billion, had a mean loss equal to the baseline estimate of \$450 million, and implied a 5 percent chance that losses would exceed \$1 billion in any year.

Having chosen the probability distribution, analysts used standard mathematical software to estimate the annual expected value of Treasury payouts under the proposal—which totaled \$96 million.¹¹ The estimate was not greatly affected by the specifics of the distribution: changing the parameters to allow the chance of losses above \$1 billion to range from 4 percent to 7 percent would have yielded estimated annual costs ranging from \$84 million to \$112 million. Even a far narrower distribution, with only a 1 percent annual chance of losses exceeding \$1 billion, would still have implied a cost of \$56 million.¹²

Accordingly, CBO informed the House and Senate Banking Committees that it estimated the Administration's proposal as costing \$100 million per year. For whatever combination of reasons, the proposal was not included in legislation considered in either chamber.

11. Losses would be expected to exceed the threshold in 36 percent of years, and the average size of those losses when they did occur would be roughly \$770 million, including the \$500 million that would continue to be paid out of the SAIF.
12. The main reason for the low sensitivity of estimated federal costs is that the proposed threshold of \$500 million is close to CBO's baseline estimate of \$450 million. A change in parameters that increases or decreases the variance (spread) of the distribution while leaving the mean at \$450 million affects the probability and expected value of losses farther down the tail of the distribution more strongly than losses near the mean.

CONCLUSION

The examples above show that the probabilistic approach is applicable to a wide range of legislative proposals involving thresholds or other asymmetries and that it may be useful even if relevant data are scarce or the key probabilities change over time. More important, the case studies demonstrate that the approach yields better estimates of the budgetary costs or savings of such one-sided bets and thus provides the Congress with more accurate information for making its decisions. Deterministic methods will continue to predominate in CBO's analyses, given that the uncertainties surrounding most of its estimates are reasonably symmetric. Nevertheless, the probabilistic approach is now established as an important element of the agency's set of estimating tools.